

Efficient Loop Navigation for Symbolic Execution

Jan Obdržálek and Marek Trtík

Masaryk University, Brno, Czech Republic
{obdrzalek, trtik}@fi.muni.cz

Abstract. Symbolic execution is a successful and very popular technique used in software verification and testing. A key limitation of symbolic execution is in dealing with code containing loops. The problem is that even a single loop can generate a huge number of different symbolic execution paths, corresponding to different number of loop iterations and taking various paths through the loop.

We introduce a technique which, given a start location above some loops and a target location anywhere below these loops, returns a feasible path between these two locations, if such a path exists. The technique infers a collection of constraint systems from the program and uses them to steer the symbolic execution towards the target. On reaching a loop it iteratively solves the appropriate constraint system to find out which path through this loop to take, or, alternatively, whether to continue below the loop. To construct the constraint systems we express the values of variables modified in a loop as functions of the number of times a given path through the loop was executed.

We have built a prototype implementation of our technique and compared it to state-of-the-art symbolic execution tools on simple programs with loops. The results show significant improvements in the running time. We found instances where our algorithm finished in seconds, whereas the other tools did not finish within an hour. Our approach also shows very good results in the case when the target location is not reachable by any feasible path.

1 Introduction

Symbolic execution has been studied since 70's [6,21]. The main idea of symbolic execution is to represent input by symbols and then to symbolically perform operations on values dependent on these inputs. I.e. symbolic execution is a generalization of concrete execution. Symbolic execution is usually used in the context of automatic test generation. With the arrival of powerful SMT (Satisfiability-Modulo-Theories) solvers, e.g. [30,27], came a generation of powerful software tools for verification and test generation [28,7,23,26].

However symbolic execution quickly reaches its limits when confronted with loops. As loops are widely used this is a significant problem. A typical situation is that reaching a particular location below a loop depends on the number of times this loop was iterated. Even worse, reaching that location may depend not only on the number of iterations, but also on what particular paths through the loop were chosen, and the order in which they were taken. Since in symbolic execution any iteration of a loop creates a new branch in the tree of symbolic executions, the size of the tree can become very large with even a single loop. Without deriving any information about the loop symbolic execution is forced to systematically explore all branches of this tree, running out of time even on small programs.

We aim to solve the following problem: Given a start location above a piece of code containing complicated loops, including loop sequences and loop nesting, and a target location anywhere in the code below, the goal is to find some feasible path between the start and target location, if such a path exists.

The idea behind our algorithm is relatively simple. Assume we perform symbolic execution and want to reach a given target location. On reaching a loop we enquire an oracle which paths through this loop, and in which order, we should execute in order to reach the target location. Following the oracle’s advice we get to our target, building path condition along the way. Only in our approach the oracle is replaced by a constraint system, which is less powerful. For each iteration it may suggest the next path to take, or to finish iterating this loop.

To build the constraint system we express the values of variables modified in a loop as functions of the number of times a given path through the loop was executed. This concept extends the simple one of counting loop iterations. Moreover, multiple counters for each path through the loop may be needed to correctly handle loop nesting. The expressed values are then ‘merged’ over all paths through a given loop. Constraint system is then created by taking branching conditions of the code below the loop and replacing all variables with the corresponding functions of loop counters.

We suggest that our algorithm is most useful when integrated into existing tools based on symbolic execution. Since the algorithm is itself based on symbolic execution the integration should not be a big problem. Our algorithm would work as a specific search strategy, activated when a global search strategy needs to navigate to a specific program location below some complicated loop structure. This would greatly improve the loop handling ability of such symbolic execution tool and allow it to explore more code in less time.

To evaluate our approach, we have built an experimental implementation of our technique – a tool called CBA. We tested CBA on nine benchmarks we designed to capture those loop structures which often appear in practice. CBA was able to solve all the benchmarks in seconds, confirming that the approach we decided to use is correct. We also compare the performance of CBA to successful symbolic execution tools PEX [28,24] and KLEE [7] and show that, on our set of benchmarks, CBA is several orders of magnitude faster than either of these tools.

The rest of the paper is organized as follows. In Section 2 we recall the basics of symbolic execution and then show how our technique works on an example. The algorithm is explained in Section 3. In Section 4 we first describe the set of benchmarks used for testing loop handling capabilities and compare the running times of CBA to those of other tools. We also evaluate the performance data from various stages of our algorithm. We survey related work in Section 5, and conclude with Section 6.

2 Overview

In this section we start with an example showing the limitations of symbolic execution when it comes to dealing with loops. We then present an overview of our approach to solving this problem on this example. Detailed description of our algorithm is deferred to the next section.

2.1 Symbolic Execution and Loops

The idea behind symbolic execution can be explained as follows: Instead of executing the program on a concrete input, we introduce for each input variable i a symbolic value α_i , standing for some concrete, but yet unknown value from the domain of i . The execution of the program then proceeds in much the same way as normal (concrete) execution. The important difference is that now many variables can contain symbolic values. For example let a and b be input variables with symbolic values α_a and α_b . Then after executing the statement $c = 2 * a + b$ the variable c will contain the symbolic value $2 \cdot \alpha_a + \alpha_b$.

Branching is treated in the following way. During symbolic execution we maintain a boolean formula φ called the *path condition*, originally set to `true`. Assume that the symbolic execution

reached a branching statement and let ψ be the associated condition. If $\varphi \wedge \psi$ is satisfiable we continue with the true branch and we update the path condition to $\varphi \wedge \psi$. (Similarly for the condition $\varphi \wedge \neg\psi$ and false branch.) If both $\varphi \wedge \psi$ and $\varphi \wedge \neg\psi$ are satisfiable we fork the execution and work on the resulting branches independently. When the symbolic execution is terminated, e.g. by reaching the desired location, we use an SMT solver to derive concrete input values from the actual path condition.

Symbolic execution works very well on programs with complicated branching sequences [8,28,24]. Using a SMT solver symbolic execution can effectively generate inputs systematically examining all branches of the code. However, symbolic execution reaches its limits when confronted with a code containing loops. Presence of even a single simple loop in a program may cause a very large (or even infinite) number of forks in the symbolic execution.

The problem with loops can be demonstrated on the program in Figure 1 (a). The goal is to find a feasible path to the `assert` statement on line 9. It is easy to see (at least for a human) that more than one such path exists and that it must iterate both the loops. However, there are about 2^{30} different execution paths which must be explored in the worst case to show that the code at line 9 is reachable. The problem here is that the condition refers to the values of `a` and `b`, which depend on the input (the arrays `A` and `B`) only indirectly. Thus the condition on line 8 does not affect the path condition of any trace and therefore it is not possible for a SMT solver to compute an input leading to the `assert` branch. Even though this seems to be a very simple example, it took the symbolic execution tool PEX 99 seconds to find an input reaching the `assert` statement. Moreover when we substituted the predicate `a>12` on line 8 with `a>17` (thus line 9 becomes unreachable), PEX was not able to finish within 5 hours.

1 <code>int a=0, b=0;</code>	c₀	c₁	c₃
2 <code>for (int i=0; i<15; ++i)</code>	<code>a=0</code>	<code>i<15</code>	<code>j<15</code>
3 <code> if (A[i]==1)</code>	<code>b=0</code>	<code>A[i]==1</code>	<code>B[j]==2</code>
4 <code> ++a;</code>	<code>i=0</code>	<code>++a</code>	<code>++b</code>
5 <code> for (int j=0; j<15; ++j)</code>	<code>i>=15 : {c₁, c₂}</code>	<code>++i</code>	<code>++j</code>
6 <code> if (B[j]==2)</code>	<code>j=0</code>		
7 <code> ++b;</code>	<code>j>=15 : {c₃, c₄}</code>	c₂	c₄
8 <code> if (a>12 && a+b==23)</code>	<code>a>12</code>	<code>i<15</code>	<code>j<15</code>
9 <code> assert(0);</code>	<code>a+b==23</code>	<code>A[i]!=1</code>	<code>B[j]!=2</code>
		<code>++i</code>	<code>++j</code>
(a)		(b)	

Fig. 1. Example used throughout Section 2. (a) C program containing loops. (b) Its chain program form.

2.2 Algorithm Overview

We introduce our technique of handling loops on the example above, in three distinct phases:

Phase 1: Conversion to chain normal form To better facilitate reasoning about loops we represent the program using linear code fragments called *chains*. The decomposition of our program to chains (what we call *chain program form* later in the text) is shown in Figure 1 (b). Chain c_0 is the topmost chain (called *root chain* later in the paper), corresponding to a path through the code where we replace the outermost loops by constructs of the form $\varphi : \{c_1, c_2, \dots\}$ with the following meaning: at this point chains c_1, c_2, \dots may be executed any number of times and in any order, but the condition φ must hold after we finish executing them. Note that the condition on line 8 was replaced by a pair of assertions.

As to the other chains, chain c_1 represents the path through the loop on lines 2–4 which goes through the positive branch of the `if` statement and c_2 the only other path through this loop. The same holds for the chains c_3 and c_4 and the loop at lines 5–7. One can easily see that there is a natural correspondence between the program (Figure 1 (a)) and its linear representation (Figure 1 (b)).

The task of finding some feasible path to the `assert` statement now depends on finding a proper interleaving of chains c_1 and c_2 for the first loop, and c_3 and c_4 for the second one.

$\begin{array}{l} \mathbf{c_1} \\ i(\kappa_1) = \kappa_1 + \alpha_i \\ a(\kappa_1) = \kappa_1 + \alpha_a \\ \{\mathbf{c_1}, \mathbf{c_2}\} \\ i(\kappa_1, \kappa_2) = \kappa_1 + \kappa_2 + \alpha_i \\ a(\kappa_1) = \kappa_1 + \alpha_a \end{array}$	$\begin{array}{l} \mathbf{c_2} \\ i(\kappa_2) = \kappa_2 + \alpha_i \\ a(\kappa_2) = \alpha_a \end{array}$	$\begin{array}{l} \mathbf{c_3} \quad \mathbf{c_4} \\ j(\kappa_3) = \kappa_3 + \alpha_j \quad j(\kappa_4) = \kappa_4 + \alpha_j \\ b(\kappa_3) = \kappa_3 + \alpha_b \quad b(\kappa_4) = \alpha_b \\ \{\mathbf{c_3}, \mathbf{c_4}\} \\ j(\kappa_3, \kappa_4) = \kappa_3 + \kappa_4 + \alpha_j \\ b(\kappa_3) = \kappa_3 + \alpha_b \end{array}$
--	--	--

$\begin{array}{ll} (1) & \kappa_1 + \kappa_2 \geq 15 \\ (2) & \kappa_1 + \kappa_2 - 1 < 15 \quad \text{if } \kappa_1 + \kappa_2 > 0 \\ (3) & \kappa_3 + \kappa_4 \geq 15 \end{array}$	$\begin{array}{ll} (4) & \kappa_3 + \kappa_4 - 1 < 15 \quad \text{if } \kappa_3 + \kappa_4 > 0 \\ (5) & \kappa_1 > 12 \\ (6) & \kappa_1 + \kappa_3 = 23 \quad \kappa_1, \kappa_2, \kappa_3, \kappa_4 \in \mathbb{N} \end{array}$
---	--

Fig. 2. Top: Recurrent variables expressed as functions of counters, including the functions after merging.

Bottom: Constraint system $\mathcal{S}(c_0)$ of the root chain c_0 .

Phase 2: Building a constraint system We start by expressing the values of variables in each chain (except root chains) as functions of the number of times this chain was executed – κ_i . Each chain c_i is linked to *chain counter* κ_i , which takes values from \mathbb{N}_0 . The link is given by the bottom index of the counter. We show how to compute the values of variables on chain c_1 , using counter κ_1 . Let α_i and α_a be the initial symbolic values of variables `i` and `a`, which are not known to this chain. Then $i(\kappa_1) = \kappa_1 + \alpha_i$ and $a(\kappa_1) = \kappa_1 + \alpha_a$ are the values of these variables expressed as functions of κ_1 . The functions for the other chains are shown in Figure 2 (top). In terms of the original program we have introduced a counter for each unique path through each loop.

Now for any given variable `i` and each path through a given loop there may be different function expressing the value of `i` in terms of the relevant counter. In the second step we try to express the value of `i` by a single function of multiple counters. This abstracts from any concrete interleaving of the subchains, but the value of the variable is expressed precisely. So in the case of chains c_1 and c_2 the value of `i` can be expressed as $i(\kappa_1, \kappa_2) = \kappa_1 + \kappa_2 + \alpha_i$. The results for our example are presented in Figure 2 under the headings $\{\mathbf{c_1}, \mathbf{c_2}\}$ and $\{\mathbf{c_3}, \mathbf{c_4}\}$.

We can now build a constraint system for the topmost chain c_0 . The constraints are obtained by processing all its assertions. There are four assertions in the chain c_0 : `i>=15`, `j>=15`, `a>12`, and `a+b==23`. We replace the variables by their previously computed values (i.e. functions of counters), arriving the constraint system $\mathcal{S}(c_0)$ depicted at Figure 2 (bottom). The constraints (1), and (2) came from the assertion `i>=15`, (3), and (4) from the assertion `j>=15`, (5) from `a>12`, and finally (6) from `a+b==23`. The constraint (1) was computed as follows. First we substitute variables in the assertion by their values, obtaining $i(\kappa_1, \kappa_2) = \kappa_1 + \kappa_2 + \alpha_i \geq 15$. α_i represents the value of `i` on reaching the `i>=15` : $\{c_1, c_2\}$ instruction. Here $\alpha_i = 0$, giving us the constraint (1), which speaks about the values of κ_1 and κ_2 just after the associated loop was executed for the last time. However, this also means that for all previous executions, where the values are $\kappa'_1 \leq \kappa_1$ and $\kappa'_2 \leq \kappa_2$ such that $\kappa'_1 + \kappa'_2 < \kappa_1 + \kappa_2$, the negated condition $i(\kappa'_1, \kappa'_2) < 15$ must hold – i.e. there is an additional constraint for each such choice of κ'_1 and κ'_2 . This can be rephrased as

$\kappa_1 + \kappa_2 - a < 15$ for $a \in \{1, 2, \dots, \kappa_1 + \kappa_2 - 1\}$. Our experimentation shows that it is sufficient to take only a single constraint for $a = 1$, giving us the constraint (2). Constraints (3), (5) and (6) are derived similarly to (1) and constraint (4) in the same way as (2). Note that we do not construct constraint systems for chains c_1, c_2, c_3, c_4 since they do not contain any subchains.

The point of the constructed constraint system $\mathcal{S}(c_0)$ is that only those executions which reach the `assert` statement satisfy $\mathcal{S}(c_0)$ (if we instantiate the counters by the number of times the corresponding path through a loop was executed). Which in turn means that solving our constraint system will limit the space of counter values we need to consider. E.g. from (5) we know that $\kappa_1 \in [13, \infty]$, and therefore the chain c_1 , which is linked to the counter κ_1 , must be executed at least 13 times on any feasible path. Similarly once we know the value for κ_1 , then from the constraint (6) and the fact that κ_1 is not modified in the second loop (there is no link between κ_1 and either c_3 or c_4) we can derive the number of times chain c_3 needs to be executed.

Phase 3: Navigating the symbolic execution With the chains, counters and constraint systems in place we may proceed with the final stage of the algorithm – finding some feasible path to line 9. We do this by employing slightly modified symbolic execution. We initialize all counters to 0 and proceed down the chain c_0 in a standard way until we reach the line 4: `i>=15`: $\{c_1, c_2\}$ (i.e. the entry point of the first loop). There are two subchains c_1 and c_2 for this loop, linked to counters κ_1 and κ_2 . Now we iteratively do the following:

- Check whether we can improve current solution of the system by incrementing κ_1 or κ_2 . If we cannot, we stop iterating and continue down the chain c_0 .
- Otherwise we call a decision procedure to tell us which counter to increment. This procedure will be described in more detail in Section 3.
- Lets assume κ_1 was chosen. In that case we symbolically execute the chain linked to κ_1 , i.e. c_1 . We also increment the counter κ_1 .

Note that it may happen that the symbolic execution may get stuck because a wrong choice of counter (counters) to increment. In that case we backtrack, asking the decision procedure for the next best counter to be incremented. Also note that by first establishing the values of counters κ_1 and κ_2 , before proceeding any further in the chain c_0 , we significantly cut down the number of paths which need to be explored.

Having solved the loop related to chains c_1 and c_2 we proceed with the execution, handling the loop related to chains c_3 and c_4 in the same way. Once we arrive at the end of c_0 we return the current path condition, which identifies a feasible path. In our case one such path condition is $A[0] = 1 \wedge \dots \wedge A[12] = 1 \wedge A[13] \neq 1 \wedge A[14] \neq 1 \wedge B[0] = 2 \wedge \dots \wedge B[9] = 2 \wedge B[10] \neq 2 \wedge \dots \wedge B[14] \neq 2$. By construction this path condition is also the sought-for path condition for the original program.

2.3 Additional Notes

Our technique holds a significant edge over the standard symbolic execution in proving no feasible path to a target location exists. Consider the example in Figure 1 where the predicate `a>12` on line 8 is replaced by `a>17`. The only change in the constraint system is (5) to (5') $\kappa_1 > 17$. From the constraints (2), and (5') we can derive that κ_2 has to have a negative value. Since all counters obtain values from \mathbb{N}_0 , the constraint system has no solution. This means there is no feasible path to the target location. The important fact is that we were able to prove this even before starting the symbolic execution.

The existence of a solution to the constraint system does not guarantee existence of a feasible path. There are two possible problems: 1) not all solutions correspond to feasible paths that lead

to the target location, and 2) single solution may correspond to multiple paths with different interleaving of the iterations through a loop. So while the constraint system effectively prunes the paths which do not lead to the target, it is not able to give us a feasible path on its own. We therefore still need to use the symbolic execution. For the same reasons each time we fail we need to backtrack to check the other possible solutions.

3 Algorithm

As we have explained in the previous section, the algorithm proceeds in a three phases. In this section we give detailed description of all three phases. Before we present our algorithm we want to state its limitations: Currently our technique is intraprocedural – i.e. we do not support function calls. We also deal only with integer variables and arrays, and handle neither heap manipulation operations nor pointer arithmetic.

3.1 Phase 1: Programs as Chains

In this section we describe how to convert a program to chain program form. It may be helpful for the reader to follow the example in Fig. 1. We understand a program P to be an oriented graph, in which the vertices are the program instructions and edges express the control flow. We write $u \rightarrow v$ ($u \rightarrow^* v$) if there is an edge (path) from u to v . We assume that there is a single start and a single terminal vertex (s_0 and t_0) and that the program P contains no unreachable code. Moreover we assume that the successors of a branching vertex are labeled c and $!c$ (for some condition c), indicating what condition must hold in order to enter the corresponding branch. Converting e.g. a C program to this form is obvious (program in our definition is basically a control flow graph). We also require the program to be in the static single assignment (SSA) form, i.e. for each program variable there is at most one place this variable is assigned to (using the standard conversion).

We define the *chain program form* $C(P)$ of P to be the set of all chains in P . A *chain* in P is a path in P which is of one of the two specific types: *Root chain* is a simple path (no vertex appear twice) $s_0 \rightarrow^* t_0$. *Subchain* is a simple path $v' \rightarrow^* v$ such that it is a suffix of some path $\pi : s_0 \rightarrow^* v \rightarrow v' \rightarrow^* v$ in P where v is the only vertex which appears twice in π . (If there are two different paths $s_0 \rightarrow^* v$, then the same path $v' \rightarrow^* v$ is treated as two different subchains.) In the rest of the paper we treat chains as linear sequences of vertices, and call their vertices *nodes*. In our example c_0 is the root chain, and $c_1 \dots c_4$ are the subchains. The set of all chains can be easily obtained by unfolding the graph of P into a tree, starting in s_0 and stopping each time a vertex is repeated on a path from s_0 (or when we reach t_0).

In chains there are three types of nodes – assume nodes, transform nodes and loop nodes. *Assume nodes*, e.g. $a > 12$ in c_0 , correspond to branching conditions. *Transform nodes*, e.g. $j = 0$ in c_0 , correspond to assignment statements which change the programs state. Finally *loop nodes*, e.g. $i \geq 15 : \{c_1, c_2\}$ in c_0 , are those nodes, from which there is at least one edge in P to the first vertex of some subchain. We call such a subchain a chain *associated* to this node (c_1 and c_2 in this case). Note that each subchain corresponds to a unique path through a loop and there can be many subchains associated to the same loop node. Moreover each subchain can also contain loop nodes, each having its own associated subchains. In the following two phases of the algorithm we assume that there is only one root chain. If there are multiple root chains, we run the remaining two phases of the algorithm separately for each root chain. The results can then be combined in an obvious way.

Let $C(P)$ be the chain program form associated to a program P . Then an *execution path* in $C(P)$ is a sequence of nodes, which is created as follows: we take some root chain and take the

nodes one by one. On reaching a loop node, we may either continue with the next node in the chain, or choose one of the subchains associated with this loop node. In that case we take this subchain and proceed recursively. On reaching the end in the subchain we go “one level up” to the associated loop node in the parent chain and repeat our choice to either take the next node of the parent chain or choose another associated subchain. We finish once we reach the terminal node for the root chain. It is easy to check that the following statement holds:

Theorem 1 *The algorithm described above converts each program P to chain normal form $C(P)$ such that for each path in P there is a corresponding execution path in $C(P)$ and vice versa. (By correspondence we mean that the sequences of instructions along these two paths are the same). Moreover if P is in SSA form, then so is each chain of $C(P)$.*

Exponential growth of chain program form It can be seen that representing a program P in chain program form can bring an exponential blowup in size. Such blowup is caused by the presence of branching statements. However, in our experience the number of chains is quite often very low. On the other hand it is not difficult to come up with a program for which the transformation to chain program form will actually cause an exponential increase in size. Indeed, such a growth can be observed for three of our benchmarks Hello/HW/HWM in Section 4.2.

However, compared to vanilla symbolic execution our approach still offers significant improvements. If we look at a symbolic execution tree of even a very simple loop structure, we can see that every path through a loop (represented by a single chain in our case) can appear many times in the tree – both on the same branch and on different branches. In other words, the size of the chain program form is usually much smaller than the standard symbolic execution tree for the same loop structure. Moreover, in the last stage of the algorithm (symbolic execution of $C(P)$) the use of constraint systems allows us to early prune branches not leading to the terminal node. So we can have significant space and time savings over vanilla symbolic execution despite the exponential growth of chain program form.

3.2 Phase 2: Building the Constraint Systems

In our approach constraint systems are used to guide the symbolic execution in search for a feasible path from the start to the target node. Here we show how to build the constraint system $S(c)$ for each chain c . An important idea behind the construction is to express the values of variables used in loops as functions of counters for the subchains. The counters in each constraint system are linked to concrete chains. This link between constraint systems, counters and concrete chains (in the chain program form) is the key idea of the algorithm.

The pseudocode for this phase is shown as Algorithm 1. We proceed using modified symbolic execution. The modification is twofold: First, it works on chains, not programs. Second, the domain of symbolic values is extended to contain counters (and expressions using counters) and a special value \star with the intended meaning “do not know”. (Any expression containing \star evaluates to \star .)

Let us take a chain c . At the beginning each variable i has a symbolic value α_i and the constraint system $S(c)$ is empty. Next we symbolically execute the chain: Handling of the transform nodes is clear. Assume nodes are treated as sources of constraints for $S(c)$. Each assertion is first instantiated with the current values of variables, and then inserted to the constraint system only if it references some counter. (As per the example in Section 2.2 multiple constraints can be produced from a single assertion.) On reaching a loop node n we first recursively build the constraint systems for all subchains associated to this node, obtaining symbolic values of variables (which can now depend on counters of some (possibly nested) subchains). For each variable we then merge the symbolic

input : chain $c = i_1, i_2, \dots, i_k$ (i_j is the j -th instruction)
output: constraint system $\mathcal{S}(c)$, and symbolic values of variables
 $\mathcal{S}(c) = \emptyset$
for $j=0; j < k; j++$ **do**
 switch node type of i_j **do**
 case Loop node
 for $l=1; l < |\text{subchains}(n)|; l++$ **do**
 BuildConstraintSystem(l -th subchain)
 merge symbolic values returned from subchains
 update the symbolic state of c with merged values
 for each variable v s.t. c is the reset chain for v **do**
 remember that c is the reset chain for v
 case Assume node
 instantiate the assertion a (using the current symbolic state)
 if a contains a counter **then**
 add relevant constraints to the constraint system $\mathcal{S}(c)$ of c
 break
 case Transform node
 modify the current symbolic state according to the associated assignment statement
for each variable v in c **do** express the symbolic value as a function of the temporary counter κ_c^v
return symbolic values of variables

Algorithm 1: BuildConstraintSystem

values obtained in the subchains (see the section *Merging values ...* below). The current symbolic state of c is then updated with the merged values. At this point we also detect the variables for which this chain is the reset chain (see the section *Expressing values ...* below). Since each loop node has an associated branching condition, we finish processing this condition as we would for the assume node. Finally, when we reach the end of the chain, we express the values of variables as functions of loop counters (and return these values). $\mathcal{S}(c)$ now contains the complete constraint system for the chain c . We now describe the process in more detail. We start by explaining the last step, because it is here where counters are dealt with and the notion of recurrent variables introduced.

Expressing values using counters Let us fix a chain c . The goal is for each variable to compute a function expressing its value in terms of counters. We focus on so called *recurrent variables*, which are the variables whose value 1) changes on the execution path corresponding to c and 2) their value is function of their initial value before executing c . An example of a recurrent variable is the variable i in the chain c_1 , for which we get $i = \alpha_i + 1$. To detect recurrent variables for a given chain c we simply analyze symbolic state resulting from symbolic execution of this chain.

For each recurrent variable we express its value in terms of how many times the chain c was executed – using the counter κ_c associated with the chain c . In our example, $i(\kappa_1) = \alpha_i + \kappa_1$. We use a very simple custom difference equation solver to express the values of variables using counters. Our solver handles only those recurrences which correspond to arithmetic (e.g. $i = \alpha_i + 7$) and geometric (e.g. $i = 3 \cdot \alpha_i$) progressions. This restriction is justified by the fact that, according to our experience, overwhelming majority of code uses only such progressions. In case we are not able to solve a recurrence, we use the “do not know” value \star .

An important point to make is that the initial value for i , α_i , can be set by some chain r , of which the current chain c is a subchain. Therefore the value of i does not depend only on the number of times c was executed, but, more specifically, on the number of times c was executed since last execution of r . Therefore the value of i in fact depends on a counter κ_c^r parametrized by two chains: the *update chain* c and the *reset chain* r – i.e. $i(\kappa_c^r) = \alpha_i + 2 \cdot \kappa_c^r$. This counter is

incremented each time the chain c is executed, and set to zero each time the chain r executed. If there is no reset chain for a given variable, we use the plain counter κ_c , where c is the update chain and the root chain is used as the reset chain. Note that all counters used in our example in Figure 1 are of this type. The following statement is true for chain program forms, and follows from the fact all chains are in the SSA form:

Lemma 1 *Let v be a recurrent variable whose update chain is c . Then v is not reset (to its initial value) in any subchain of c and there exists at most one superchain of c where v is reset.*

At the time of processing the update chain c for v we do not know yet what superchain of c is the reset chain for v . Therefore we use a temporary counter κ_c^v to express the value of v . When processing a chain d such that 1) d is a superchain of c and 2) the value of v is no longer given by a recurrence expression (α_v does not occur in the symbolic value of v), we know that d is the reset chain for v . We remember this information, and once all constraint systems are built we replace all occurrences of each temporary counter κ_c^v in all constraint systems with the correct counter κ_c^d .

Merging values from subchains We explain the merging process on the case of two subchains. The extension to multiple subchains is straightforward.

Let us assume that a chain has two subchains with the associated counters being κ_c and κ_d (as the reset chains are not important here, we omit the upper indices) and there is a variable i value of which is expressed as $i = i_1(\kappa_c)$ in the first chain and $i = i_2(\kappa_d)$ in the second. We would like to “merge” the values of i – i.e. to find a function $i(\cdot, \cdot)$ such that $i = i(\kappa_c, \kappa_d)$. Let α_i be the symbolic value of i on entering the subchains. There are some simple cases: e.g. if $i_1(\kappa_c) = i_2(\kappa_d) = v$ for some constant v , then obviously also $i(\kappa_c, \kappa_d) = v$. Similarly if $i_1(\kappa_c) = i_2(\kappa_d) = \alpha_i$. On the other hand if $i_1(\kappa_c) = v_1 \neq v_2 = i_2(\kappa_d)$ then there is no such function $i(\kappa_c, \kappa_d)$. In that case we put $i(\kappa_c, \kappa_d) = \star$.

The most interesting case is when both $i_1(\kappa_c)$ and $i_2(\kappa_d)$ depend on α_i – e.g. $i_1(\kappa_c) = v_1 \cdot \kappa_c + \alpha_i$ and $i_2(\kappa_d) = v_2 \cdot \kappa_d + \alpha_i$. This means that the value of i is updated in both subchains. In this case we can easily derive that $i(\kappa_c, \kappa_d) = v_1 \cdot \kappa_c + v_2 \cdot \kappa_d + \alpha_i$. Table 1 sums up all supported merge operations for functions of a single counter. In all other cases we put $i(\kappa_c, \kappa_d) = \star$. From Table 1 we see that our ability to merge is limited. E.g. we are not able to merge even $i_1(\kappa_c) = v_1 \cdot \kappa_c + \alpha_i$ and $i_2(\kappa_d) = \alpha_i \cdot v_2^{\kappa_d}$. Also \star propagates quickly through the system and constraints with \star are not useful in our approach. On the other hand we can merge functions with different numbers of counters, e.g. $i(\kappa_c, \kappa_d)$ with $i(\kappa_3)$ etc.

$i_1(\kappa_c)$	$i_2(\kappa_d)$	$i(\kappa_c, \kappa_d)$
α_i	α_i	α_i
v	v	v
$v_1 \cdot \kappa_c + \alpha_i$	$v_2 \cdot \kappa_d + \alpha_i$	$v_1 \cdot \kappa_c + v_2 \cdot \kappa_d + \alpha_i$
$\alpha_i \cdot v_1^{\kappa_c}$	$\alpha_i \cdot v_2^{\kappa_d}$	$\alpha_i \cdot v_1^{\kappa_c} \cdot v_2^{\kappa_d}$

Table 1. Supported merge operations for unary functions

3.3 Phase 3: Constraints-Driven Symbolic Execution

The last stage of our algorithm is to navigate (modified) symbolic execution in order to find a feasible path from s_0 to t_0 . We modify standard symbolic execution in order to run on the chain program form described in Section 3.1. To do so, we first extend the symbolic state by extra variables representing the values of counters. Second, on entering a chain, we instantiate all symbols α_v in the constraint system associated with the chain by their actual symbolic values.

The symbolic execution starts by setting all counters for which the current chain is the reset chain to zero and then proceeds on the root chain as normal until it reaches a loop node, which will play a role of a branching statement in standard symbolic execution. A symbolic execution tool typically asks an oracle (a heuristic), when an execution reaches forking branch. Since two or more branches can be simultaneously taken from that point, an oracle is responsible to choose a branch which is more likely to reach the goal of exploration than others. In our case branching points are the loop nodes, and the oracle is the decision procedure given by Algorithm 2.

input : c, D :: a chain and an subset of its sub-chains
 A :: constraint system for c
output: Chosen chain (i.e. c or some $d \in D$) or **null**.
if A has no solution **then return null**
if counters' values represent a solution of A **then return** c
 $R := \{\text{reset chains of counters for (subchains of) } D, \text{ whose reset gets closer to a solution of } A\}$
 $U := \{\text{update chains of counters for (subchains of) } D, \text{ whose update gets closer to a solution of } A\}$
if $R \cup U = \emptyset$ **then return** c
else return arbitrary element from $R \cup U$

Algorithm 2: chooseChain

Let c be the currently executed chain, A its (instantiated) constraint system, i the processed loop node, and D be the subset of the set of subchains associated to i (containing those subchains which have not been yet explored during backtracking). If A has no solution, we immediately stop symbolic execution for this branch. Otherwise, if the current values of counters already form a solution of A , we continue executing c , as there is no reason to execute any of the subchains. Otherwise we need to choose a chain $d \in D$ which, hopefully, brings us closer to a solution of A . If there is such d , we continue with the symbolic execution of d . Finally if there is no such d , then we also continue executing c , hoping that we can closer to a solution of A at some loop node below.

Now we describe what we mean by “getting closer to a solution of A ”. Let w be a vector of current values of all the counters such that w is not a solution to A . We now ask whether there is a vector v on natural numbers such that 1) $v + w$ is a solution to A , and 2) there is a counter κ such that $d \in D$ (or some of its subchains) is the update chain for κ (reset chain for κ) and there is a positive (negative) number in the corresponding position in v . If yes, then executing the chain d gets us “closer to a solution of A ”.

There are many possible approaches to compute the vector v . One is to simply use a SMT solver to obtain a solution to A . In our implementation we use interval abstraction: we overapproximate the set of all solutions by giving a set of intervals for each counter. These intervals are derived from the constraint system, and we have a solution to A if the value of each counter lies in one of its intervals. In this abstraction individual components of any solution vector are independent. Thus choosing some vector v is trivial.

Finally we have to say what happens when the symbolic execution reaches the terminal node of a chain c . We first increment all the associated counters κ_c^d (for all d). If c is a subchain we continue by (again) executing the associated loop node in the parent chain, otherwise c is a root chain and we reached the target node.

We conclude by stating the soundness and incompleteness of our method (the latter follows immediately from incompleteness of the standard symbolic execution):

Theorem 2 (Soundness) *If the symbolic execution of $C(P)$ (as described in Section 3) terminates with success, then the returned path condition represents a feasible path from start to target instruc-*

tion in the original program P . Moreover if the symbolic execution fails, then there is no feasible path in P to the target instruction.

Theorem 3 (Incompleteness) *There exists a program P with reachable target instruction for which the symbolic execution of $C(P)$ never terminates.*

4 Experimental Results

To evaluate the effectiveness of our technique we implemented it (with all the restrictions mentioned at the beginning of Section 3) in our tool CBA, and tested it on a set of nine benchmarks. We also compared the performance of CBA to that of two very successful tools PEX [28,24] and KLEE [7]. All the nine benchmarks share some common properties: 1. the code contains loops (so the benchmarks produce a huge symbolic execution tree) 2. there is a unique location to be reached 3. they consist of only one function (since our technique does not handle function calls). In the first six benchmarks the goal is to find a feasible path to the target location. On the other hand in the last three benchmarks there is no feasible path to the target location and the goal is to show that no feasible path exists.

Benchmark Description The first three benchmarks **Hello/HW/HWM** are adapted from [1] (there is only verbal description, no code). The HWM benchmark accepts a C string as an input and scans the string for the presence of substrings "Hello", "World", "At" and "Microsoft!". HW and Hello are simplified versions of the HWM benchmark, looking for the first two words (one word) only.

In **DOIF** we model a typical piece of code which scans an input and, for each member of the input array, performs an action which depends on its value. This benchmark is supposed to exercise primarily the third stage of the algorithm. Branching inside the loops enormously expands the number of paths in the model. **DOIFex** is an extension of this benchmark, and tests behaviour on sequences of loops with internal branching.

The **EQCNT** benchmark contains nested loops with branching, where a variable defined in the outermost scope is modified in the innermost loop. **EQCNTex** is a modified benchmark (in a sense two instances of EQCNT in sequence), however the number loop iterations is now given explicitly (in contrast to the two remaining benchmarks, where it is dependent on the input). For an algorithm to be efficient on this benchmark it has to aggressively prune infeasible paths.

The **OneLoop** benchmark consists of simple loop in which the variable i , with initial value 0, is increased by 4 in every iteration. Once the loop is finished we check whether $i=15$, which is false for any value of the input variable n . **TwoLoops** is an extension of the previous benchmark by adding a second loop, whose loop condition depends on the value computed in the first loop.

4.1 Tool Comparison

In this section we present the experimental results we obtained by running PEX, KLEE and our tool CBA on our set of benchmarks. We ran our test on an Intel i7/920 2.67GHz Windows machine with 6GB of RAM. Since KLEE is native C++ Linux application, we used the Cygwin library to run KLEE on Windows, resulting in an overhead caused by calls to Cygwin's dynamic library. We decided to reduce this negative effect by using the `time` utility to measure the 'user' time of KLEE. However this was as close as we could get to running all the tools in the same environment.

For PEX we present two results for each of the benchmarks. This is because the performance of PEX is affected by a set of configurable parameters. The first result is obtained in the way recommended by PEX developers – with all parameters set to infinity. The second result (indicated

by an asterisk), which is usually better, is obtained by iteratively running PEX and adjusting the parameters according to suggestions provided with the unsuccessful runs.

Comparison results The results are presented in Table 2. We measured the time required to reach the target location. Each benchmark has an associated timeout (column **timeout**), which was set according to the perceived difficulty of that particular benchmark. The success was defined as reaching the target location (or demonstrating it is not possible to reach this location) within the specified time limit.

Looking at the table one can see that, on our set of benchmarks, CBA significantly outperforms both PEX and KLEE. This shows that on short pieces of code containing non-trivial loops our technique can effectively guide the symbolic execution to the chosen target location. On the other hand PEX and KLEE clearly suffer from the limitations of the symbolic execution when dealing with loops.

Test	timeout	PEX	PEX *	KLEE	CBA
Hello	30m	3.234s	7.233s	0.093s	0.026s
HW	1h	14.890s	11.107s	37m 0s	0.175s
HWM	1h	fail	8m 54s	timeout	1.997s
DOIF	30m	timeout	20m 27s	timeout	0.388s
DOIFex	1h	timeout	timeout	timeout	1.745s
EQCNT	30m	1m 43s	11.592s	timeout	0.191s
EQCNTex	1h	46m 12s	42m 20s	timeout	2.458s
OneLoop	30m	2m 14s	4m 27s	timeout	0.002s
TwoLoops	30m	1m 4s	57.426s	timeout	0.003s

Table 2. Running times of PEX, KLEE and CBA.

4.2 Performance Analysis of CBA

In this section we discuss the behaviour of CBA on our set of nine benchmarks. Table 3 shows the performance data. The three enclosing columns refer to the three stages of our algorithm: *Chain prog. form* refers to the conversion of a program into chain program form. *Chains* gives the number of root/all chains, *Time* the time needed for the conversion and *Space* the size of resulting data structures. *Constr. Systems* covers the second stage. *Elim* shows how many root chains were shown to be infeasible even before getting to the last stage and *Size* gives the total number of constraints left after pruning. Finally *Constraints-Driven Sym. Exe.* refers to the last stage. *SStat* gives the number of symbolic states (i.e. vertices of a symbolic execution tree) visited. *CSOL* gives the number of calls to the constraint solver: The first number is for the initial solution, the second one for the remaining calls. *SMT* is the number of calls to the Z3 SMT solver and finally *PC* counts the number of predicates in the resulting path condition.

The number of chains for the HWM benchmark is quite large. There were 161 chains, including 81 root chains. The high number of chains for HWM is reflected in the time and space needed to build the chain program form. We can see the negative effect of exponential growth of chain program form here. If we compare the running time and the number of chains for the three related benchmarks Hello/HW/HWM we see that number of chains grow indeed exponentially.

Another interesting observation is that it is hard to predict how long will the last stage take based on the performance of the first two stages. To see this, consider the results obtained for the HWM and DOIFex benchmarks. In the case of HWM there are 161 chains (before pruning) and 20 constraints in remaining chains (after pruning), while for DOIFex there are only 9 chains and 4 constraints. However in both cases the last stage explores a comparable number of symbolic states

Test	Chain Prog. Form			Constr. Systems				Constraints-Driven Sym. Exe.				
	Chains	Time	Space	Elim	Size	Time	Space	SStat	CSOL	SMT	Time	PC
Hello	3/6	0.003s	1kB	2	5	0.001s	3.1kB	10	8 / 36	19	0.023s	5
HW	9/17	0.009s	20 kB	8	10	0.040s	6 kB	44	30 / 170	92	0.144s	10
HWM	81/161	0.097s	369 kB	80	20	0.719s	12 kB	174	112 / 686	376	1.456s	22
DOIF	1/5	0.003s	1 kB	0	3	0.014s	2 kB	98	97 / 349	136	0.380s	26
DOIFex	1/9	0.004s	3 kB	0	4	0.008s	5 kB	211	209 / 728	212	1.757s	26
EQCNT	1/4	0.004s	1 kB	0	3	0.003s	2 kB	45	44 / 245	45	0.187s	43
EQCNTex	1/7	0.004s	2 kB	0	6	0.005s	4 kB	1192	1022 / 7286	1233	2.458s	0
OneLoop	1/2	0.003s	293 B	0	2	0.001s	698 B	0	1 / 0	0	0.001s	0
TwoLoops	1/3	0.003s	578 B	0	2	0.001s	1 kB	0	1 / 0	0	0.001s	0

Table 3. Performance data for CBA on our set of benchmarks.

in similarly comparable time. This indicates that, in respect to our algorithm, the number of chains and constraints are not the only important parameters.

The last negative output can be seen on the EQCNTex benchmark. The values for the last stage are an order of a magnitude higher than for the other tests. Note that this is despite the number of chains being very low. Remember that in EQCNTex there is no feasible path to the target location, and many paths need to be explored to prove it. EQCNTex benchmark shows the limitations of our algorithm with respect to solving such problems. Even though it can effectively prune away many paths (as witnessed by the last two benchmarks, OneLoop and TwoLoops) this is not always sufficient. Nevertheless CBA still fared significantly better on EQCNTex than both PEX and KLEE.

5 Related Work

The earliest work dealing with symbolic execution [6,21] showed that symbolic execution can be an effective approach to test generation. However the astronomical blowup of program model caused by loops was not in the centre of interest. Usability evaluation of symbolic execution for proving correctness of program implementing Floyd’s method [10] was in [21], but problems with loops were handled by manually inserting ASSUME statements where necessary.

Modern effective techniques based on symbolic execution are mostly hybrid, combine symbolic execution with some other approaches. The first group are techniques based of combining (alternating) concrete and symbolic execution [13,26,28,15]. This approach primarily avoids the problems caused by limitations of SMT solvers. Although the practical usability is greatly improved, these techniques have no effect on the ability to handle loops. The second group combines symbolic execution with some validation technique [16,19,2,23,17]. This approach is much more successful from the point of handling loops. Thanks to employing the complementary techniques, many symbolic paths can be effectively pruned away when exploring the symbolic state space. This can often lead to effective navigation of symbolic execution in programs with loops. There is also a group of techniques which aim to make symbolic execution effective in the general case, not specifically focused on just programs with loops [5,11,1,8,7,14].

The idea of using constraint system for analyzing loops was considered before in different contexts. First approach, dating back to 70’s, infers relations between program variables [20,9], while the more recent techniques are primarily focused on formal verification, and inductive invariant computation [3,18]. Analysis of loops using loop-counters as the artificial program variables is also well known [22].

The technique of Loop-Extended Symbolic Execution [25] (LESE) is probably the one most closely related to our approach. The LESE approach introduces symbolic variables for the number of times each loop was executed, and links these with features of a known input grammar such as

variable-length or repeating fields. This allows the symbolic constraints to cover a class of paths that includes different number of loop iterations, expressing loop-dependent program values in terms of properties of the input.

Our approach is very different: Instead of extending the input by new symbolic variables to reason about multiple symbolic execution paths at once, our goal is to build a constraint systems to steer the symbolic execution through loops towards a specified target. For this reason we introduce counters which are linked to different paths through a cycle, contrasting to the overall iteration count used by the LESE approach. Our technique therefore applies to a much more general class of programs.

Finally there is an orthogonal line of research which tries to improve the symbolic execution for programs with some special types of inputs. Some examples are techniques for dealing with programs with string inputs [4,29], and techniques which reduce input space given by an input grammar [12,25]. These approaches can be effective on loops when such loops are closely related to the input.

6 Conclusion and Future Work

In this paper we introduced a new algorithm for effective navigation of symbolic execution through loop containing code. The algorithm infers a collection of constraint systems and uses them to steer the symbolic execution towards a target location. To build these constraint systems we express the values of variables modified in a loop as functions of the number of times a particular path through the loop was executed.

We have also built an experimental implementation of our technique and tested its effectiveness on a set of nine benchmarks. Our tool was able to correctly solve each of these benchmarks within seconds, being several orders of magnitude faster than the leading symbolic execution tools. Moreover we have demonstrated that our technique is also useful for proving that no feasible path to a target location exists.

Finally, we argue that it would be beneficial for general-purpose tools based on symbolic execution to integrate our technique as a new search strategy. This strategy would then be activated each time the symbolic execution needs to navigate to a specific target location below some complicated loop structure. Since our algorithm is itself based on symbolic execution, such integration should not be too difficult.

There are many interesting open directions for future work. An obvious task would be to extend our approach to interprocedural setting. Moreover, so far we have considered only integer variables and arrays. It would be interesting to extend our technique to handle more sequential containers (e.g. lists or vectors) and/or floating point arithmetics. Another approach is to try to curb the growth of the chain program form, for example by merging those chains which have the same effect on program execution. Finally it would be nice to actually integrate our approach with the existing symbolic execution tools like KLEE.

Acknowledgements We would like to thank Nikolai Tillmann for promptly answering our questions regarding PEX. We also thank Václav Brožek, Vojtěch Forejt, Antonín Kučera and Jan Strejček for their helpful comments on earlier drafts of this paper.

References

1. S. Anand, P. Godefroid, and N. Tillmann. Demand-driven compositional symbolic execution. In *TACAS'08*, volume 4963 of *LNCS*, pages 367–381. Springer, 2008.

2. N. E. Beckman, A. V. Nori, S. K. Rajamani, R. J. Simmons, S. Tetali, and A. V. Thakur. Proofs from tests. In *ISSTA '08*, pages 3–14. ACM, 2008.
3. D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Path invariants. In *PLDI '07*, pages 300–309. ACM, 2007.
4. N. Bjørner, N. Tillmann, and A. Voronkov. Path feasibility analysis for string-manipulating programs. In *TACAS '09*, pages 307–321. Springer-Verlag, 2009.
5. P. Boonstoppel, C. Cadar, and D. Engler. RWset: attacking path explosion in constraint-based test generation. In *TACAS'08/ETAPS'08*, pages 351–366. Springer-Verlag, 2008.
6. R. S. Boyer, B. Elspas, and K. N. Levitt. SELECT—a formal system for testing and debugging programs by symbolic execution. In *Proceedings of the international conference on Reliable software*, pages 234–245. ACM, 1975.
7. C. Cadar, D. Dunbar, and D. Engler. KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *OSDI'08*, pages 209–224. USENIX Association, 2008.
8. C. Cadar, V. Ganesh, P. M. Pawlowski, D. L. Dill, and D. R. Engler. EXE: Automatically generating inputs of death. *ACM Trans. Inf. Syst. Secur.*, 12(2):1–38, 2008.
9. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *POPL '78*, pages 84–96. ACM, 1978.
10. R. W. Floyd. Assigning meanings to programs. In *Proceedings of a Symposium on Applied Mathematics*, volume 19, pages 19–31, 1967.
11. P. Godefroid. Compositional dynamic test generation. In *POPL '07*, pages 47–54. ACM, 2007.
12. P. Godefroid, A. Kiezun, and M. Y. Levin. Grammar-based whitebox fuzzing. In *PLDI '08*, pages 206–215. ACM, 2008.
13. P. Godefroid, N. Klarlund, and K. Sen. DART: directed automated random testing. In *PLDI '05*, pages 213–223. ACM, 2005.
14. P. Godefroid, M. Y. Levin, and D. A. Molnar. Active property checking. In *EMSOFT '08*, pages 207–216. ACM, 2008.
15. P. Godefroid, M. Y. Levin, and D. A. Molnar. Automated whitebox fuzz testing. In *Network Distributed Security Symposium (NDSS)*, pages 151–166, 2008.
16. P. Godefroid, A. V. Nori, S. K. Rajamani, and S. D. Tetali. Compositional may-must program analysis: unleashing the power of alternation. In *POPL '10*, pages 43–56. ACM, 2010.
17. B. S. Gulavani, T. A. Henzinger, Y. Kannan, A. V. Nori, and S. K. Rajamani. SYNERGY: a new algorithm for property checking. In *SIGSOFT '06/FSE-14*, pages 117–127. ACM, 2006.
18. S. Gulwani, S. Srivastava, and R. Venkatesan. Program analysis as constraint solving. In *PLDI '08*, pages 281–292. ACM, 2008.
19. A. Gupta, R. Majumdar, and A. Rybalchenko. From tests to proofs. In *TACAS '09*, pages 262–276. Springer-Verlag, 2009.
20. M. Karr. Affine relationships among variables of a program. *Acta Informatica*, 6:133–151, 1976.
21. J. C. King. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, 1976.
22. F. Nielson, H. I. Nielson, and Ch. Hankin. *Principles of Program Analysis*. Springer, 2005.
23. A. V. Nori, S. K. Rajamani, S. Tetali, and A. V. Thakur. The Yogi project: Software property checking via static analysis and testing. In *TACAS '09*, pages 178–181. Springer-Verlag, 2009.
24. <http://research.microsoft.com/Pex>.
25. P. Saxena, P. Poosankam, S. McCamant, and D. Song. Loop-extended symbolic execution on binary programs. In *ISSTA '09*, pages 225–236. ACM, 2009.
26. K. Sen, D. Marinov, and G. Agha. CUTE: a concolic unit testing engine for C. In *ESEC/FSE-13*, pages 263–272. ACM, 2005.
27. <http://sites.google.com/site/stpfastprover/>.
28. N. Tillmann and J. de Halleux. Pex – white box test generation for .NET. In *TAP'08*, volume 4966 of *LNCS*, pages 134–153. Springer, 2008.
29. R. G. Xu, P. Godefroid, and R. Majumdar. Testing for buffer overflows with length abstraction. In *ISSTA '08*, pages 27–38. ACM, 2008.
30. <http://research.microsoft.com/projects/Z3>.